PATENT

Case Docket No. MERCURY.054A
Date: April 16, 2003

In re application of	:	Eli Levy
Appl. No.	:	09/531,821
Filed	:	March 21, 2000
For	:	SERVER MONITORING USING VIRTUAL POINTS OF PRESENCE
Examiner	:	Quang N. Nguyen
Art Unit	:	2141

I hereby certify that this correspondence and all marked attachments are being deposited with the United States Postal Service as first-class mail in an envelope addressed to: Board of Patent Appeals and Interferences, United States Patent and Trademark Office, P.O. Box 2327 Washington, D.C. 20231, on

April 16, 2003
(Date)

Ronald J. Schoenbaum, Reg. No. 38,297

BOARD OF PATENT APPEALS AND INTERFERENCES
UNITED STATES PATENT AND TRADEMARK OFFICE
P.O. Box 2327
WASHINGTON, D.C. 20231

Sir:

Transmitted herewith in triplicate is an Appellants' Brief to the Board of Patent Appeals:

(X) Fee for filing brief in the amount of $320.00; and
(X) A return prepaid postcard.

If applicant has not requested a sufficient extension of time and/or has not paid any other fee in a sufficient amount to prevent the abandonment of this application, please consider this as a Request for an Extension for the required time period and/or authorization to charge our Deposit Account No. 11-1410 for any fee which may be due. Please credit any overpayment to Deposit Account No. 11-1410.

Ronald J. Schoenbaum
Registration No. 38,297
Attorney of Record
Customer No. 20,995
(949) 760-0404

H:\DOCS\ROS\ROS-5300.DOC\dns
041403

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant	:	Eli Levy	Group Art Unit: 2141
Appl. No.	:	09/531,821	
Filed	:	March 21, 2000	
For	:	SERVER MONITORING USING VIRTUAL POINTS OF PRESENCE	
Examiner	:	Quang N. Nguyen	

APPELLANTS' BRIEF
PURSUANT TO 37 C.F.R. § 1.192

Board of Patent Appeals and Interferences
Washington, D.C. 20231

Dear Sir:

Appellant, Applicant in the above-captioned patent application, appeals the final rejection set forth in the Final Office Action mailed on March 11, 2003. A check for the filing fee is enclosed. Please charge any additional fees that may be required now or in the future to Deposit Account No. 11-1410.

I. REAL PARTY IN INTEREST

The real party of interest in the present application is Mercury Interactive Corporation.

II. RELATED APPEALS AND INTERFERENCES

No related appeals or interferences are pending.

III. STATUS OF CLAIMS

Claims 1-40, all of which were finally rejected by the Examiner, are currently pending in the application and are the subject of this appeal. The claims, as amended in response to the first Office Action, are attached hereto as an appendix.

IV. STATUS OF AMENDMENTS

No amendments were made in response to the Final Office Action.

V. SUMMARY OF THE INVENTION

As described in the background section of the present application, the performance of a web site or other Internet server system, as experienced by end users of the system, can vary significantly depending on the geographic locations of these users. For example, users in London may experience much greater response times than users in San Francisco. Such variations in end user response times may occur, for example, as the result of Internet traffic conditions, malfunctioning Internet routers, or malfunctioning DNS (Domain Name Service) servers. Because these types of problems can often be corrected once identified, is it desirable to provide a mechanism for allowing web site operators to monitor the performance of their web sites as seen from various geographic locations. Present application at page 1, line 19 to page 2, line 2.

The prior art addresses this problem by providing a web site monitoring system in which agent computers are set up at selected Internet connection points, typically in major population centers such as major cities throughout the world. Each such agent computer runs special monitoring software (also referred to as agent software) that causes the agent computer to periodically access the target web site from its respective geographic location as a simulated web site user, and to monitor the web site's response times. The agent computers thus measure the end-user performance (performance as seen by end users) of the target web site as seen from their respective geographic locations. The agent computers report the resulting performance data (response time measurements and possibly other performance parameters) over the Internet to a centralized database, where the data is typically aggregated and made available to the web site operators for viewing. The agent computers and centralized database are typically operated by a monitoring service provider. Present application at page 2, lines 3-14.

A significant problem with this prior art approach is that the cost of setting up and maintaining agent computers in many different geographic regions is typically very high. For example, the monitoring service provider typically must pay for regional personnel who have been trained to set up and service the agent software and computers. The monitoring service provider may also incur costs for maintaining the security of the agent computers, and for upgrading the agent software as new versions become available. Present application at page 2, lines 17-23.

The present invention overcomes the above and other problems by setting up the agents computers in one or more centralized locations or "data centers," rather than deploying an agent computer at each of the geographic locations for which end-user performance is to be monitored. These centrally located agent computers are connected to remote Internet connection points that reside in the various geographic locations for which end user performance is to be monitored. For example, a data center residing in Sunnyvale California may be connected to Internet connection points in Seattle, New York, Dallas, London, and various other cities to permit monitoring of end-user performance of web sites as seen from these cities. See, e.g., the "Summary of the Invention" section and page 4, lines 24 to page 6, line 5 of the present application. These geographically distributed Internet connection points are referred to as "virtual points of presence" or VPOPs, as they serve as virtual agent locations on the Internet.

A significant benefit of this approach is that the agent software resides on computers in the one or more central data centers, rather than on computers in each of the various geographic locations for which end user performance is to be monitored. Thus, there is no need to deploy and maintain agent software or other monitoring components within the each of the various geographic locations for which end user performance is to be monitored. The cost of setting up and maintaining the monitoring system therefore is significantly reduced. See present application at page 3, lines 5-7 and page 5, line 28 to page 6, line 5.

In Applicant's preferred embodiment, request messages (HTTP requests, etc.) generated by the centrally located agent computers are transmitted over communications links to the desired VPOPs, where these messages flow onto the Internet and then to the target server system. These requests appear to the target server system to emanate from users who are local to the VPOPs. Thus, for example, a single data center may generate message traffic for accessing and

monitoring a particular web site from New York, San Francisco, London, and various other cities. Responses from the target server system flow back from the VPOPs to the centrally located agent computers over the same communications links used for the corresponding requests.

The links between the agent computers and the VPOPs are preferably special communications links, such as dedicated connection-oriented links, for which latency is known or can be reliably estimated. This allows the delay attributable to a given agent-VPOP link to be deducted from the overall response time seen by the agent computer, so that the response time as seen from the VPOP location can be accurately measured. For instance, if an agent computer in Sunnyvale California accesses a target web site via a VPOP in New York, the delay attributable to the Sunnyvale-New York communications link can be subtracted from the overall response time seen by the agent computer. In one embodiment, the agent-VPOP links include dedicated links for which the link delays can be estimated based on the current loads on such links. In another embodiment, the delays caused by the agent-VPOP links are sufficiently small to be ignored. Present application at, e.g., page 3, lines 8-22.

VI. ISSUES PRESENTED ON APPEAL

The following issue is presented: whether Claims 1-40 are properly rejected under 35 U.S.C. § 103 as unpatentable over Chen (U.S. Patent 5,812,780) in view of Boss (U.S. Patent 6,157,618), collectively referred to herein as the "applied references."

VII. GROUPING OF CLAIMS

All of the rejected claims in the present application should not stand or fall together. Appellants, for purposes of this appeal only, are grouping the claims as follows:

GROUP 1: Claims 37-40. The sole independent claim of this group is Claim 37, which is directed generally to a system for monitoring a web site from multiple geographic locations. The system includes a data center comprising at least one computer system programmed to access and monitor the web site. The system further includes a plurality of Internet connection points, each of which is located at a different respective one of the multiple geographic locations, with at least some of the Internet connection points being geographically remote from others and

from the data center. The data center sends request messages to the web site, and receives response messages from the web site, through the plurality of Internet connection points to monitor the web site as seen by users in each of the multiple geographic locations. Group 1 also includes dependent Claims 38-40, which further define the invention of Claim 37.

GROUP 2: Claims 12-20, and 22. The sole independent claim of Group 2 is Claim 12, which is directed to a system for monitoring the performance of a server system as seen from multiple user locations. The remaining claims depend from Claim 12. Unlike the broadest claim of Group 1, the claims in Group 2 include language requiring that the monitoring system monitor the performance of the server system as seen from the multiple user locations "without the need for performance monitoring components local to said user locations."

GROUP 3: Claims 30 and 32-36. The sole independent of this group is Claim 30, which is directed to a system for monitoring access to an Internet server system on a public computer network. The remaining claims depend from Claim 30. Unlike the broadest claims in Groups 1 and 2, the claims in Group 3 call for "a dedicated communications link that connects the agent computer to an Internet point of presence that is geographically remote from the agent computer." In addition, the Group 3 claims call for an agent computer that "is configured to use IP (Internet Protocol) addresses associated with the Internet point of presence, such that all forward and reverse message traffic between the agent computer and the Internet server system flows across the dedicated communications link."

GROUP 4: Claims 1, 4-11, 23, and 26-29. The sole independent claims of this group are Claims 1 and 23, which are directed to the associated methods of the invention. Claims 3-11 depend from Claim 1, and Claims 25-29 depend from Claim 23.

GROUP 5: Claims 2, 21, 24 and 31. This group consists of dependent claims that include language related to subtracting or deducting from the total response time the delay associated with the link to the remote access location or VPOP.

GROUP 6: Claims 3 and 25. This group consists of dependent claims that further define corresponding claims of Group 5 by requiring that the delay associated with the link be estimated based at least in part on the load on that link.

Although additional patentable distinctions are recited in many of the other dependent claims, Appellant wishes to limit the issues on appeal by grouping the claims as indicated above.

VIII. ARGUMENT

A. THE CLAIM GROUPS ARE PATENTABLY SEPARATE

As set forth above, the claims of Group 1 call for a data center that sends request messages to the web site, and receives response messages from the web site, through a plurality of Internet connection points at different respective geographic locations to monitor the web site as seen by users in each of the multiple geographic locations. As discussed below, this aspect of the Group 1 claims provides a patentable distinction over the applied references.

The claims of Group 2 require that the monitoring system monitor the performance of the server system as seen from the multiple user locations "without the need for performance monitoring components local to said user locations." This aspect of the Group 2 claims is not disclosed or suggested by the applied references (as discussed below), and provides a separate basis for patentability over these references.

All of the claims of Group 3 call for "a dedicated communications link that connects [an] agent computer to an Internet point of presence that is geographically remote from the agent computer," and further state that the agent computer "is configured to use IP (Internet Protocol) addresses associated with the Internet point of presence, such that all forward and reverse message traffic between the agent computer and the Internet server system flows across the dedicated communications link." As discussed below, these limitations of the Group 3 claims similarly provide a separate basis for patentability over the applied references.

The claims of Group 4 are directed to the associated method of the invention, and include method limitations that provide additional bases for patentability over the applied references. For instance, independent Claim 1 recites "transmitting the first request message over a first communications link from the monitoring location to a first network access location that is remote from the monitoring location for transmission on the network at the first network access location, *to thereby simulate access to the server system by a user who is local to the first network access location.*" The other independent claim of this group–independent Claim 23–includes similar language.

The claims of Group 5 include language related to subtracting or deducting from the total response time the delay associated with the link to the remote access location. As discussed

below, this feature of the invention provides an additional basis for patentability over the applied references.

The claims of Group 6 further require that the delay associated with the link be estimated based at least in-part on a load on that link. As discussed below, this aspect of the invention provides an additional basis for patentability over the applied references.

In view of the foregoing, the Appellant submits that the six claim groups are patentably separate.

B. DISCUSSION OF REFERENCES RELIED UPON BY EXAMINER

1. Chen

The Chen patent discloses a system in which a client component, referred to as LoadSim, simulates the actions of many concurrent users to apply a load to an Exchange server. To the Exchange server, the LoadSim client appears equivalent to many different Exchange clients running on separate client computers on a local area network. As the load is applied to the Exchange server, the LoadSim client monitors response times of the Exchange server.

In contrast to Applicant's preferred embodiment, and to the invention as set forth in some of Applicant's claims, the system disclosed in Chen does not monitor the performance of the server as seen from multiple geographic locations without the need to have performance monitoring components in such geographic locations. In this regard, the LoadSim client appears to monitor the performance of the target server only as seen from the computer on which the LoadSim client runs, and not as seen from other geographic locations.

In addition, no mechanism is disclosed or suggested in Chen for allowing the LoadSim client to monitor the server's performance as seen from locations that are geographically remote from this LoadSim client computer. In this regard, Chen's objective is apparently to measure the load capacity of the Exchange server, and not to evaluate performance of the Exchange server as seen from multiple different geographic locations.

2. Boss

Boss discloses a system in which a number of geographically distributed "data gathering" client computers access a target site to measure the performance of the site as seen from their respective geographic locations. The data gathering functionality of the client computers is

provided by a data-gathering client component 402 that is executed by each data gathering computer. These data gathering computers report their performance measurements to a central "UseMon" server for analysis. The UseMon server also informs the data gathering computers of the identities of the target servers to be accessed/monitored. The client/data gathering computers may be computers of ordinary Internet users, who may download and install the data gathering component on their respective computers. These client computers may access the Internet via dial-up connections to ISPs.

As with the system disclosed in Chen, the system of Boss does not monitor the performance of the target server as seen from the multiple geographic locations without the need for performance monitoring components in such geographic locations. Rather, in the system of Boss, the data-gathering client (a performance monitoring component) resides and executes on a computer in each geographic location for which end user performance is monitored. For example, to monitor a web site's response times as seen by users in the New York City area, it would be necessary to run the data-gathering client component 402 on a computer in the New York City area.

C. DISCUSSION OF THE ISSUES ON APPEAL

For the reasons set forth below, Appellant respectfully submits that the teachings of the Boss and Chen patents do not render the present invention obvious.

1. The applied references do not disclose or suggest every limitation of the broadest claim of Group 1

In order to establish prima facie obviousness of a claimed invention, all of the claim limitations must be taught or suggested by the prior art. See MPEP § 2143.03. As set forth below, Chen and Boss fail to satisfy this requirement with respect to Claim 37 (the broadest claim of Group 1), and thus fail to render obvious the claims of Group 1.

Claim 37 reads as follows:

> 37. A system for monitoring a web site from multiple geographic locations, the system comprising:
>
> a data center comprising at least one computer system programmed to access and monitor the web site; and

> a plurality of Internet connection points, each of which is located at a different respective one of said multiple geographic locations, at least some of the Internet connection points being geographically remote from others and from the data center;
>
> wherein the data center sends request messages to the web site, and receives response messages from the web site, through the plurality of Internet connection points to monitor the web site as seen by users in each of the multiple geographic locations.

The Examiner did not provide a separate analysis of Claim 37, but instead relied on his analysis of independent Claim 1. See Final Office Action at section 16. In rejecting Claim 1, the Examiner took the position that the UseMon server of Boss represents the "monitoring location" of Claim 1, and that the data gathering computers of Boss represent the "first network access location" and "second network access location" of Claim 1. Appellant therefore assumes that the Examiner is treating the UseMon server of Boss as the "data center" required by Claim 37, and is treating the data gathering computers of Boss as the recited "Internet connection points."

This interpretation ignores the language in Claim 37 stating that the data center "sends request messages to the web site, and receives response messages from the web site, through the plurality of Internet connection points to monitor the web site as seen by users in each of the multiple geographic locations." In this regard, the UseMon server of Boss does not either send request messages to the web site being monitored, or receive response messages from the web site being monitored, via the data gathering computers. Rather, as indicated above, the UseMon server simply notifies the data gathering computers of the details of the target web site to be monitored, and collects the resulting response time measurements generated by the data gathering computers. The UseMon server of Boss thus cannot properly be treated as the "data center" recited in Claim 37.

Chen similarly fails to disclose or suggest a data center as set forth in Claim 37. In this regard, neither the LoadSim client nor any other component of Chen "sends request messages to [a] web site, and receives response messages from the web site, through [a] plurality of Internet connection points [in different respective geographic locations] to monitor the web site as seen by users in each of the multiple geographic locations."

Because Claim 37 includes limitations that are not disclosed or suggested by the applied references, the obviousness rejection of the Group 1 claims is improper.

2. The applied references do not disclose or suggest every limitation of the broadest claim of Group 2.

Group 2 consists of independent Claim 12, and Claims 13-20 and 22 which depend from Claim 12. Claim 12 reads as follows:

> 12. A system for monitoring performance of a server system as seen from multiple user locations, including user locations that are geographically remote from one another, the system comprising:
>
> a plurality of agent computers that are locally interconnected at a central location, the agent computers configured to generate request messages to the server system as simulated users and to monitor responses from the server system to the request messages; and
>
> a plurality of communications links, each communications link connecting one or more of the plurality of agent computers to a respective network access point within a respective user location that is geographically remote from the central location;
>
> wherein the plurality of agent computers are configured to communicate with the server system via the plurality of communications links to monitor performance of the server system as seen from the multiple user locations;
>
> whereby the system monitors the performance of the server system as seen from multiple user locations without the need for performance monitoring components local to said user locations.

Appellant respectfully submits that Boss and Chen do not disclose or suggest all of the limitations of Claim 12. For example, Boss and Chen do not disclose or suggest a system that "monitors the performance of [a] server system as seen from multiple user locations without the need for performance monitoring components local to said user locations." As discussed above, neither Chen nor Boss discloses or suggests such a capability. Rather, the systems disclosed in Chen and Boss appear to measure the performance (response times) of the target server system only as seen from the location or locations of the client/agent computers that execute the associated performance monitoring software. Specifically, Chen's system measures server response times as seen from the client computer that runs the LoadSim client; and Boss's system measures server response times as seen from each data gathering computer. Neither reference discloses or suggests a mechanism that would allow a server's performance, as seen from multiple user locations, to be monitored without the need to install and run monitoring software in such user locations.

Boss and Chen also fail to disclose or suggest the particular configuration set forth in the first three subparagraphs of Claim 12. In connection with these subparagraphs, the Examiner appears to be treating the data gathering computers of Boss as the "agent computers," and

treating the dial-up connections between these computers and their respective ISPs as the recited "communication links." See Final Office Action at page 8, last paragraph. The data gathering computers of Boss, however, do not appear to be "locally interconnected at a central location" as recited in Claim 12. In addition, the dial-up connections of Boss do not connect the data gathering computers to access points in "respective user location[s] that [are] geographically remote from the central location" (note that each circle shown in Figure 13 of Boss represents a "local geographic area," and that the data gathering computers are shown as being connected only to ISPs within their respective local geographic areas). Thus, the first three subparagraphs of Claim 12 cannot properly be read on the disclosure of Boss.

Because Boss and Chen fail to disclose or suggest all of the limitations of Claim 12, the obviousness rejection of the Group 2 claims is improper.

3. The applied references do not disclose or suggest every limitation of the broadest claim in Group 3.

Group 3 consists of independent Claim 30, and dependent Claims 32-36 which depend from Claim 30. Claim 30 reads as follows:

> 30. A system for monitoring access to an Internet server system on a public computer network, comprising:
>
> an agent computer that is configured to access the Internet server system as at least one simulated user while monitoring performance of the Internet server system; and
>
> a dedicated communications link that connects the agent computer to an Internet point of presence that is geographically remote from the agent computer;
>
> wherein the agent computer is configured to use IP (Internet Protocol) addresses associated with the Internet point of presence, such that all forward and reverse message traffic between the agent computer and the Internet server system flows across the dedicated communications link;
>
> whereby the agent computer monitors performance of the Internet server system as seen from the geographically remote Internet point of presence.

Boss and Chen do not disclose all of the limitations of Claim 30. For example, neither reference discloses or suggests "a dedicated communications link that connects the agent computer to an Internet point of presence that is geographically remote from the agent computer." In this regard, neither the communications paths between the UseMon server and the data gathering computers of Boss, nor the dial-up connections between the data gathering computers and their ISPs, represent such a dedicated link.

In addition, Chen and Boss do not disclose or suggest an agent computer that "is configured to use IP (Internet Protocol) addresses associated with [a geographically remote] Internet point of presence, such that all forward and reverse message traffic between the agent computer and the Internet server system flows across the dedicated communications link," in the context of the other claim limitations. As discussed in the present application, this feature provides a mechanism for allowing an agent computer to communicate with the target server system via–and to monitor the server system's response times as seen from–the geographically-remote Internet point of presence. (See present application at, e.g., page 3, lines 12-15; page 5, lines 25-27; and page 6, lines 19-22.) In connection with this claim language, the Examiner relies on the disclosure at column 3, lines 49-56 and column 4, lines 55-58 of Boss. Final Office Action at page 9, first two paragraphs. These portions of Boss do not, however, disclose or suggest this feature of Claim 30.

Because Boss and Chen do not disclose all of the limitations of independent Claim 30, the obviousness rejection of the Group 3 claims is improper.

4. The applied references do not disclose or suggest every limitation of any claim in Group 4.

Group 4 consist of independent method Claims 1 and 23, and certain dependent claims that depend from Claims 1 and 23. As set forth below, both of these independent claims include limitations that are not disclosed or suggested by the applied references. The rejection of the Group 4 claims is therefore improper.

Independent Claim 1

Claim 1 calls for a system in which user access to a server system as experienced by users local to first and second network access locations "is monitored without the need for monitoring components local to the first and second network access locations." As discussed above, neither Chen nor Boss discloses or suggests such a capability. Rather, the systems disclosed in Chen and Boss appear to measure the performance (response times) of the target server system only as seen from the location or locations of the client/agent computers that execute the associated performance monitoring software. Specifically, Chen's system measures server response times as seen from the client computer that runs the LoadSim client; and Boss's system measures server response times as seen from each data gathering computer. Neither reference discloses or

suggests a mechanism for measuring server performance as seen from a particular geographic location without the need to run special monitoring software in that geographic location.

With respect to the other limitations recited in Claim 1, and namely those recited in the first four subparagraphs of the claim, the Examiner attempts to construe these limitations so as to cover the system of Boss. In doing so, the Examiner takes the position that these claim limitations cover a system in which a computer at the recited "monitoring location" (interpreted as the UseMon server of Boss) does not actually communicate with the target server system, but rather instructs other computers (i.e., the data gathering computers) to access and monitor the target server. See Final Office Action beginning at page 2, section 4. Appellant respectfully submits that this interpretation is improper. In this regard, the phrase "first and second request messages that represent requests from users of the server system" cannot reasonably be construed to cover transmissions from the UseMon server to the data gathering computers of information identifying the target site to be accessed.

Independent Claim 23

Claim 23 recites a method in which "the performance of the server system as experienced from the network access location is measured without a need for any performance monitoring components at the network access location." As discussed above with respect to Claim 1, Chen and Boss do not disclose or suggest this feature.

Chen and Boss also fail to disclose or suggest the method set forth in the first four subparagraphs of the claim.

Because Boss and Chen do not disclose all of the limitations recited in either independent Claim 1 or independent Claim 23, the obviousness rejection of the Group 4 claims is improper.

5. The applied references do not disclose or suggest the limitations added by the claims in Group 5.

Group 5 consists of dependent Claims 2, 21, 24 and 31. These claims are patentable in view of their dependencies from specific claims discussed above, namely independent Claims 1, 12, 23, and 30 (respectively). In addition, all of these claims recite additional limitations related to subtracting or deducting from the total response time the delay associated with the link to the remote access location or VPOP. Specifically:

Claim 2 calls for "determining a total response time to the first request message as observed at the monitoring location; and subtracting from the total response time a delay associated with the first communications link, to thereby determine a response time as experienced at the first network access point;"

Claim 21 calls for agent computers that are "configured to measure response times as seen from the user locations by measuring total response times to request messages, and by deducting round-trip latencies associated with the communications links;"

Claim 24 calls for "determining the response time as experienced at the network access location comprises determining a total response time observed at the first location, and subtracting from the total response time a round-trip delay attributable to the communications link;" and

Claim 31 calls for and agent computer or component that is "configured to measure a total response time to a request message sent to the Internet server system, and to deduct from the total response time a round-trip delay associated with the dedicated communications link, to thereby determine a response time as seen from the Internet point of presence."

In connection with the delay subtracting/deducting feature set forth in these claims, the Examiner appears to rely solely on the disclosure at column 6, lines 15-19 of Boss. See Final Office Action at sections 5 and 13-15. This portion of Boss does not, however, disclose or suggest the feature set forth in these claims. In this regard, the generation of a "full-page download time," as disclosed in Boss, does not suggest or imply that the delay associated with a communications link is subtracted or deducted from the total response time.

The rejection of the Group 5 claims is therefore improper.

6. The applied references do not disclose or suggest the limitations added by the claims in Group 6.

Group 6 consists of dependent Claims 3 and 25. These claims are patentable in view of their dependencies from specific Group 5 claims discussed above, namely Claims 2 and 24. In addition, Claims 3 and 25 recite additional limitations related to estimating the delay associated with the communications link. Specifically:

Claim 3 recites "estimating the delay based at least on a load applied to the first communications link;" and

Claim 25 recites "estimating the round-trip delay based at least upon a load applied to the communications link."

In connection with the limitations added by these claims, the Examiner again appears to rely solely on the disclosure at column 6, lines 15-19 of Boss. See Final Office Action at sections 5 and 14. This portion of Boss does not, however, disclose or suggest estimating the delay associated with the communications link. Indeed, there would apparently be no reason to make such an estimation in the systems of Chen and Boss. The rejection of the Group 6 claims is therefore improper.

7. The Examiner has not identified a valid suggestion for combining the applied references, and no such suggestion exists in the references.

Appellant further respectfully submits that in rejecting the claims of Groups 1-6, the Examiner has failed to identify a sufficient suggestion or motivation to combine or modify Boss or Chen. Appellants further submit that no such suggestion or motivation exists within these references.

In connection with this issue, Chen appears to be concerned with testing the load capacity of a target server, and not with measuring performance as seen from various geographic locations as in Boss. Because the load capacity of a server system ordinarily would not be dependent upon the user locations from which the server system is accessed, there would be little, if any, reason for one skilled in the art to incorporate the distributed monitoring teachings of Boss into the load testing system of Chen. In fact, doing so would apparently be contrary to Chen's objective of reducing the number of client computers used to load test the target server system.

Similarly, since Boss is not concerned with testing the load capacity of the target server system, there would be little, if any, reason to incorporate the teachings Chen into the data gathering clients of Boss.

Thus, in addition to the reasons set forth above with respect to specific claim groups, the rejection of the claims in Groups 1-6 is improper because no valid basis has been set forth for combining the teachings of Boss and Chen.

IX. CONCLUSION

For the reasons set forth above, Appellant submits that the obviousness rejection of Claims 1-40 is improper, and requests that the rejection be reversed.

Respectfully submitted,

KNOBBE, MARTENS, OLSON & BEAR, LLP

Dated: 4-16-03

By: _____

Ronald J. Schoenbaum
Reg. No. 38,297
Attorney of Record
2040 Main Street
Irvine, CA 92614
949-721-2950

APPENDIX – PENDING CLAIMS

1. A method of monitoring access to a server system as experienced from multiple user locations on a computer network, comprising:

at a monitoring location, generating first and second request messages that represent requests from users of the server system;

transmitting the first request message over a first communications link from the monitoring location to a first network access location that is remote from the monitoring location for transmission on the network at the first network access location, to thereby simulate access to the server system by a user who is local to the first network access location;

transmitting the second request message over a second communications link from the monitoring location to a second network access location that is remote from both the monitoring location and the first network access location for transmission on the network at the second network access location, to thereby simulate access to the server system by a user who is local to the second network access location; and

at the monitoring location, monitoring responses received from the server system to the first and second request messages;

whereby user access to the server system as experienced by users local to the first and second network access locations is monitored without the need for monitoring components local to the first and second network access locations.

2. The method as in Claim 1, wherein monitoring responses received from the server system comprises:

determining a total response time to the first request message as observed at the monitoring location; and

subtracting from the total response time a delay associated with the first communications link, to thereby determine a response time as experienced at the first network access point.

3. The method as in Claim 2, further comprising estimating the delay based at least on a load applied to the first communications link.

4. The method as in Claim 2, wherein the network is the Internet, and the first and second communications links are connection-oriented links that directly connect the monitoring location to the first and second network access locations, respectively.

5. The method as in Claim 4, wherein the first and second communications links are Asynchronous Transfer Mode links.

6. The method as in Claim 4, wherein the server system comprises a web server of a publicly accessible web site.

7. The method as in Claim 1, wherein the first and second communications links are connection-oriented links that directly interconnect the data center with the first and second network access locations, respectively.

8. The method as in Claim 7, wherein the first and second communications links are Asynchronous Transfer Mode links.

9. The method as in Claim 8, wherein the network is the Internet.

10. The method as in Claim 1, wherein transmitting the first request message over a first communications link comprises transferring the first request message over the first communications link to a router that is directly connected to the network.

11. The method as in Claim 1, wherein transmitting the first request message over a first communications link comprises transferring the first request message over the first communications link to a modem that is locally connected to the network.

12. A system for monitoring performance of a server system as seen from multiple user locations, including user locations that are geographically remote from one another, the system comprising:

a plurality of agent computers that are locally interconnected at a central location, the agent computers configured to generate request messages to the server system as simulated users and to monitor responses from the server system to the request messages; and

a plurality of communications links, each communications link connecting one or more of the plurality of agent computers to a respective network access point within a respective user location that is geographically remote from the central location;

wherein the plurality of agent computers are configured to communicate with the server system via the plurality of communications links to monitor performance of the server system as seen from the multiple user locations;

whereby the system monitors the performance of the server system as seen from multiple user locations without the need for performance monitoring components local to said user locations.

13. The system as in Claim 12, wherein each communications link is a dedicated, connection-oriented communications link.

14. The system as in Claim 12, wherein at least some of the plurality of communications links are Asynchronous Transfer Mode links.

15. The system as in Claim 12, wherein at least some of the plurality of communications links produce a round-trip latency that is predicable based on a current load on the communications link.

16. The system as in Claim 12, wherein the server system is an Internet server system, and at least some of the communications links are directly peered to the Internet within respective user locations.

17. The system as in Claim 12, wherein the plurality of agent computers are arranged into groups, wherein each group is assigned to a respective communications link and corresponding user location.

18. The system as in Claim 17, wherein the agent computers within a group are configured to use unique IP (Internet Protocol) addresses associated with a corresponding remote Internet point of presence.

19. The system as in Claim 17, wherein at least one of the groups comprises multiple agent computers that are grouped through a hub and a single port of a local switch.

20. The system as in Claim 19, wherein the switch in connected to a router that is directly connected to the plurality of communications links.

21. The system as in Claim 12, wherein the agent computers are configured to measure response times as seen from the user locations by measuring total response times to request messages, and by deducting round-trip latencies associated with the communications links.

22. The system as in Claim 12, further comprising a database that locally stores server performance data generated by the plurality of agent computers.

23. A method for monitoring access to a server system as experienced from a network access location, comprising:

at a first location which is remote from the network access location, generating a user request message that is addressed to the server system;

sending the request message over a communications link to the network access location for transmission over a computer network from the network access location to

the server system, to thereby simulate access to the server system by a user who is local to the network access location;

at the first location, receiving a response to the request message from the server system over the communications link; and

determining a response time to the user request message as experienced at the network access location;

whereby the performance of the server system as experienced from the network access location is measured without a need for any performance monitoring components at the network access location.

24. The method as in Claim 23, wherein determining the response time as experienced at the network access location comprises determining a total response time observed at the first location, and subtracting from the total response time a round-trip delay attributable to the communications link.

25. The method as in Claim 24, further comprising estimating the round-trip delay based at least upon a load applied to the communications link.

26. The method as in Claim 23, wherein the computer network is the Internet, and the communications link is a connection-oriented link that directly connects the first location to the network access location.

27. The method as in Claim 26, wherein the communications link is an Asynchronous Transfer Mode link.

28. The method as in Claim 23, wherein the server system comprises a web server of a publicly accessible web site.

29. The method as in Claim 23, further comprising recording the response time, and response times measured for other user locations, within a database that is local to the first location.

30. A system for monitoring access to an Internet server system on a public computer network, comprising:

an agent computer that is configured to access the Internet server system as at least one simulated user while monitoring performance of the Internet server system; and

a dedicated communications link that connects the agent computer to an Internet point of presence that is geographically remote from the agent computer;

wherein the agent computer is configured to use IP (Internet Protocol) addresses associated with the Internet point of presence, such that all forward and reverse message traffic between the agent computer and the Internet server system flows across the dedicated communications link;

whereby the agent computer monitors performance of the Internet server system as seen from the geographically remote Internet point of presence.

31. The system as in Claim 30, wherein the agent computer, or a component locally connected to the agent computer, is configured to measure a total response time to a request message sent to the Internet server system, and to deduct from the total response time a round-trip delay associated with the dedicated communications link, to thereby determine a response time as seen from the Internet point of presence.

32. The system as in Claim 30, wherein the dedicated communications link is a connection-oriented link.

33. The system as in Claim 30, wherein the dedicated communications link is an Asynchronous Transfer Mode link.

34. The system as in Claim 33, wherein the Asynchronous Transfer Mode link is directly peered to the Internet by a router.

35. The system as in Claim 30, wherein the agent computer is one of a plurality of agent computers that are locally coupled to a plurality of dedicated communications links by a switch, and different groups of the plurality of agent computers are assigned to different remote Internet points of presence.

36. The system as in Claim 35, further comprising a database which is locally coupled to, and configured to store performance data generated by, the plurality of agent computers.

37. A system for monitoring a web site from multiple geographic locations, the system comprising:

a data center comprising at least one computer system programmed to access and monitor the web site; and

a plurality of Internet connection points, each of which is located at a different respective one of said multiple geographic locations, at least some of the Internet connection points being geographically remote from others and from the data center;

wherein the data center sends request messages to the web site, and receives response messages from the web site, through the plurality of Internet connection points to monitor the web site as seen by users in each of the multiple geographic locations.

38. The system as in Claim 37, wherein the data center is connected to the plurality of Internet connection points by dedicated communications links.

39. The system as in Claim 37, wherein the data center monitors the web site as seen by users in the multiple geographic locations without the use of any monitoring components in any of the multiple geographic locations.

40. The system as in Claim 37, wherein the data center calculates response times of the web site as seen from each of the Internet connection points.

H:\DOCS\ROS\ROS-5312.DOC
041603